

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
(Attorney Docket No. H052617.1313US0)

TITLE

PRIORITY MECHANISM FOR SCHEDULING ISOCHRONOUS AND
ASYNCHRONOUS TRANSACTIONS ON A SHARED BUS

INVENTORS

Howard M. Brown

214 Bedford Road

Pleasantville, NY 10570

Citizenship: U.S.A.

ASSIGNEE

Hewlett-Packard Development Company, L.P.

a Delaware corporation

P.O. Box 272400

Fort Collins, Colorado 80527-2400

CERTIFICATE OF MAILING 37 C.F.R. 1.10

"Express Mail" Mailing Label Number: EV330687572US

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office To Addressee" service with sufficient postage on the date indicated below and is addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Name of Person Signing Certificate: Holly Tuynman

07/25/03
Date

Holly Tuynman
Signature

APPLICATION FOR PATENT

INVENTOR: HOWARD M. BROWN

TITLE: PRIORITY MECHANISM FOR SCHEDULING ISOCHRONOUS AND
ASYNCHRONOUS TRANSACTIONS ON A SHARED BUS

SPECIFICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of co-pending United States application Serial No. 09/514,980, filed February 29, 2000, which is incorporated herein in its entirety.

BRIEF SUMMARY OF THE INVENTION

[0002] Many computer systems utilize a shared bus that carries multiple types of traffic. For example, some shared buses carry both asynchronous and isochronous traffic. A Peripheral Component Interconnect (PCI) bus is an example of a shared bus that has been used to carry both asynchronous and isochronous traffic. While asynchronous traffic can be delayed indefinitely without adversely affecting real-time data streams, computer system performance can be adversely affected when the asynchronous traffic is delayed. Further, since isochronous traffic is time dependent, it requires a certain bandwidth and worst case latency. A multimedia stream (video and audio signal) is an example of isochronous traffic. Multimedia streams utilize an isochronous transport mechanism to ensure that data is delivered at least as fast as it needs to be displayed. Further, isochronous transport mechanisms have been utilized to ensure an audio signal is synchronized with a video signal. Provided bandwidth and latency requirements for isochronous traffic are met, latency of the isochronous traffic has not had an adverse effect on real-time performance.

[0003] A number of techniques have been employed to ensure that isochronous traffic had access to a shared bus for a given amount of time within a given frame. These methods have included static allocation of bandwidth and timer/counter based mechanisms which attempted to detect when isochronous transactions were complete. These prior methods have normally guaranteed isochronous latency at the expense of asynchronous latency. That is to say, the isochronous bus requests were fulfilled without consideration of the asynchronous bus requests. As a general rule, this resulted in increased latency for the asynchronous

requests. In addition, trying to predict isochronous transaction requirements based on time averages has normally complicated system design.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0004] A better understanding of the present invention can be obtained when the following detailed description of the preferred embodiment is considered in conjunction with the following drawings, in which:

FIG. 1 is a block diagram illustrating a typical computer system S;

FIG. 2 is a block diagram illustrating multiple devices connected by a shared bus;

FIG. 3 is a block diagram depicting a central arbitrator communicating with an isochronous device and an asynchronous device;

FIG. 4 is schematic diagram of hardware which implements an embodiment of the present invention;

FIG. 5 is a flow chart illustrating one technique for scheduling bus transactions according to an embodiment of the present invention;

FIG. 6 is a flow chart illustrating another technique for scheduling bus transactions according to a different embodiment of the present invention; and

FIG. 7 is a block diagram illustrating a network interface card which includes network driver firmware (including code to schedule asynchronous and isochronous transactions) according to an embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

COMPUTER SYSTEM

[0005] Turning to Figure 1, illustrated is a block diagram of a computer system S according to an embodiment of the present invention. Processor 100 is coupled to a host bus 104. The processor 100 in the preferred embodiment is a Pentium II processor manufactured by the Intel Corporation. Alternatively, the present invention could be implemented in a multi-processor workstation or any other computer system that utilizes a shared bus to carry asynchronous and isochronous traffic. Also coupled to the host bus 104 is a chipset 108. The chipset 108, as disclosed, is an accelerated graphics port (AGP) chipset which can be a 440BX AGPset manufactured by the Intel Corporation. The chipset 108 provides a memory controller for controlling memory 110, a host-PCI bus bridge for coupling a PCI bus 116 to

the host bus 104, and an AGP connector for connecting a graphics card 106. A display 122, as shown, is coupled to the graphics card 106.

[0006] Coupled to the PCI bus 116 is a small computer system interface (SCSI) controller 124 and a network interface card (NIC) 128. The NIC 128 can provide an Ethernet connection for coupling the computer system S to a local area network (LAN). Coupled to the SCSI controller 124 is a disk subsystem 126. The SCSI controller 124 can potentially control various disk subsystems 126 which can include these types of peripheral devices: tape drives, WORM, CD-ROM, DVD, and optical storage devices.

[0007] Also coupled to the PCI bus 116 is a PCI/ISA bridge 130. As is common, the PCI/ISA bridge 130 converts data communicated from peripherals (not shown) coupled to an ISA bus 136 for utilization on the PCI bus 118. The PCI/ISA bridge 130 also provides USB ports 132 and IEEE 1394 ports 134. A variety of peripherals, that provide asynchronous and isochronous traffic, can be coupled to the USB ports 132 and the IEEE 1394 ports 134.

[0008] The Pentium II processor 100 can be replaced with different processors other than the Pentium II without detracting from the spirit of the invention. The processor 100 is capable of running any of a number of operating systems, such as Windows 98®, Windows NT®, or a Unix-based operating system.

[0009] Again, it should be understood that a wide variety of systems could be used instead of the disclosed computer system S without detracting from the spirit of the invention. Further, other current and future operating systems could also be utilized.

SCHEDULING ASYNCHRONOUS AND ISOCHRONOUS TRANSACTIONS

[0010] The disclosed technique provides for scheduling asynchronous and isochronous transactions on a shared bus and normally reduces asynchronous latency while maintaining isochronous latency below an acceptable level. While the technique addresses asynchronous and isochronous traffic, one of skill in the art will appreciate that, it can be applied to other types of traffic, when applicable. Asynchronous traffic is an example of a primary data transaction. Isochronous traffic is an example of a secondary data transaction. Asynchronous traffic involves primary data transactions in the sense that asynchronous traffic is generally given higher priority than isochronous traffic. In the same sense, isochronous traffic in

comparison to asynchronous traffic can be considered as secondary data transactions. The scheduling techniques discussed herein in the context of asynchronous and isochronous transactions are equally applicable to other types of primary and secondary data transactions.

[0011] The shared bus can include a parallel bus such as a PCI bus or a serial bus such as an Ethernet connection carrying Ethernet data, a Universal Serial Bus (USB) or an IEEE 1394 bus. Initially, an allocated bus time is split (such as through time division multiplexing) into frames of equal length. The allocated bus time is based on factors which are typically application and system specific. When a bus request is received the technique determines whether the bus request in a current frame is for an asynchronous transaction or an isochronous transaction. As described below, the determination of the type of transaction can be accomplished in a number of ways.

[0012] If an asynchronous bus request has been initiated it is processed, otherwise an isochronous bus request (if any) is processed. Bus requests for an isochronous transaction are latched or queued if received while an asynchronous transaction is currently being processed. Asynchronous transactions are given priority until a current frame time has ended. In one embodiment, at the start of a new frame any queued isochronous transactions are processed before asynchronous transactions of a current frame are given priority. In another embodiment, queued isochronous transactions are only processed at the start of a new frame if they are from two frames prior to the new frame.

[0013] Moving to Figure 2, a plurality of devices 200-210 coupled to a shared bus 212 are illustrated. The devices 200-210 could be a variety of types, such as a server or a workstation (represented by the computer system S) or any peripheral that is normally coupled to a particular type of shared bus that carries both asynchronous and isochronous traffic. The devices 200-210 transmit and receive isochronous and asynchronous data over the shared bus 212. As previously discussed, the shared bus can be either a parallel or serial bus. For example, the shared bus can be an Ethernet connection, a PCI bus, a USB or an IEEE 1394 bus.

[0014] When the shared bus is a PCI bus, two control signals (per the PCI specification) are used for arbitration (REQ* and GNT*). Each busmaster on the PCI bus is capable of asserting a bus request signal REQ* and responding to a bus grant signal GNT* (supplied by

central arbitration logic). A typical PCI central arbitrator implementation is modeled on the central arbitration control point (CACP) as defined by microchannel architecture (MCA) or an EISA bus arbitrator. EISA uses an arbitration model with three priority levels: DMA/refresh (highest level); CPU/master (middle level); other masters (lowest level). The CACP is typically a dedicated chip which carries out bus arbitration and passes control to a busmaster. The priority of the busmasters coupled to the CACP is normally programmable.

[0015] A configuration area of 256 bytes (64 registers of 32 bits) is typically provided for each PCI device. A 64-byte header is typically included within the 256-byte configuration area. The 64-byte header normally includes a class code which indicates the type of PCI device. The class code includes a basic code and a subcode. The basic code relates the basic type of PCI device. For example, the basic code may indicate the following types of devices: a mass storage controller, a network controller, video controller, multimedia unit, a memory controller or a bridge. A subcode for a multimedia unit can indicate whether the multimedia unit is a video, audio, or other unit. Likewise, other PCI devices have corresponding subcodes. If a device is to act as a busmaster, then a busmaster (BM) bit of an instruction register (located within the configuration area) must be asserted. The configuration area can be accessed (read or write) with various functions of the BIOS interrupt 1ah.

[0016] In a PCI-based system where a CACP has been implemented, the CACP carries out bus arbitration and passes control to a busmaster. The arbitration process is initiated when multiple busmasters assert their bus request lines. Each busmaster has its own bus request output line. The CACP asserts an arbitration signal to indicate an arbitration cycle. Each busmaster that is requesting the PCI bus then outputs its priority code and compares it with other busmasters' priority code. Lower priority busmasters then deactivate their priority code. At this point, the CACP asserts a bus grant to the highest priority busmaster. Each busmaster has its own bus grant input line. The highest priority busmaster (current busmaster) then deasserts its bus request signal and takes control of the PCI bus. After the current busmaster is through using the PCI bus, or at the end of its allocated time, the current busmaster relinquishes the PCI bus. If multiple busmasters are then requesting the bus, then the arbitration process is repeated again.

[0017] Turning to Figure 3, the arbitration process is further illustrated. When an asynchronous device 310 and an isochronous device 300 assert their bus request lines, the

asynchronous device 310 is given priority, for a time, as discussed further below (in conjunction with Figures 5 and 6). The central arbitrator 330 determines which device becomes the next busmaster. The type of data (asynchronous or isochronous) provided by a device can be determined, for example, by reading a device's class code or by the priority of a device requesting the shared bus. Alternatively, a system designer may designate specific bus request inputs of an arbitrator as isochronous or asynchronous. Another approach may designate a specified amount of time at the start of each transmitted frame, in an ethernet system for example, where only left over isochronous requests would contend for the bus. Figure 4 illustrates circuitry which when implemented in conjunction with a CACP type arbitrator allows isochronous traffic to gain use of the shared bus. While Figure 3 shows an isochronous device 300 and an asynchronous device 310, it should be understood that isochronous data and asynchronous data may be provided to a shared bus by a same or different device.

[0018] Moving to Figure 4, hardware which allows isochronous traffic to gain use of the shared bus is illustrated. The disclosed hardware can be implemented within a bus arbitrator of a shared bus. Isochronous requests are clocked into a latch 400 and 402 by a frame clock signal at time (t). While only four isochronous request lines are shown, the number of isochronous request lines can be scaled as dictated by the application. Initially, the latches 400 and 402 are reset by asserting a reset line (not shown). Output data of latch 400 is valid shortly after time (t), where (t) indicates a first frame clock. Output data of latch 402 is not valid until shortly after time (t+1), where (t+1) indicates a second frame clock. Shortly after the second frame clock signal (at time (t+1)), valid data exists at the outputs of latches 400 and 402.

[0019] To prevent a potential race condition when the latches 400 and 402 are clocked, a delay 408 can be introduced such that the clocking of the latch 400 is delayed from that of the latch 402. One skilled in the art will appreciate that a delay can be implemented in a variety of ways. Each of the data outputs of the latch 400 is ANDed (AND gates 404a, 404b, 404c, and 404d) with the corresponding data output of latch 402. Outputs of AND gates 404a, 404b, 404c, and 404d are coupled to an input of OR gate 406. An output of the OR gate 406 is coupled to the bus arbitrator such that when the output of OR gate 406 is asserted high the isochronous traffic is serviced at the beginning of a new frame. This occurs after a particular isochronous request is asserted high for two successive frames. Similarly, the above concept

can be extended by adding additional latches and AND gates (with inputs corresponding to the number of latches). While the circuitry in Figure 4 is intended for use with an arbitrator, it should be understood that logic for scheduling isochronous traffic according to the invention can be implemented in a variety of ways.

[0020] Turning to Figure 5, a schedule bus transaction routine 500 is illustrated. In step 502, an allocated bus time is split into frames of equal length. From step 502, control transfers to step 504 where the routine 500 determines whether the bus request is for an asynchronous transaction. If the bus request is for an asynchronous transaction in step 504, control transfers to step 508. In step 508, the asynchronous transaction is processed. If an isochronous bus request occurs while the asynchronous transaction is being processed, then the isochronous transaction is queued. From step 508 control transfers to step 510. In step 510, the routine 500 determines whether a current frame has ended. The end of a current frame, for example, can be determined using timers. The beginning and ending of a frame are usually system and application specific. When the shared bus is an Ethernet, the beginning and ending of a frame are normally associated with a particular bit pattern in the data stream. If the current frame has ended, control transfers from step 510 to step 512. In step 512, a new frame is initiated and at the start of the new frame, any queued isochronous requests are processed in a round-robin fashion. From step 512, control returns to step 504.

[0021] In step 504, if the bus request is not for an asynchronous transaction, control transfers to step 506. In step 506, an isochronous request is processed. From step 506, control transfers to step 504. Thus, asynchronous requests are processed with the highest priority during a current frame. At the start of a new frame, any queued isochronous requests are then processed in a round-robin fashion. Utilizing this approach an isochronous request is delayed for, at most, two frames. That is, an isochronous request at the beginning of a current frame is serviced by the end of a new frame.

[0022] Moving to Figure 6, a schedule bus transaction routine 600 is illustrated. In step 602, the bus time is split into frames of equal length. As previously stated, a particular allocated bus time is application specific. For example, asynchronous traffic across a PCI bus may be limited such that various other peripherals attached to the PCI bus can utilize the PCI bus. From step 602, control transfers to step 604. In step 604, the routine 600 determines whether a bus request is for an asynchronous transaction. If the bus request is for an

asynchronous transaction, then control transfers from step 604 to step 608. In step 608, the asynchronous transaction is processed. From step 608, control transfers to step 610. In step 610, the routine 600 determines whether a current frame has ended. As above, the end of an Ethernet frame is usually detected by a particular bit pattern in the transmitted information. If the current frame has not ended, then control transfers from step 610 to step 604. If the current frame has ended, then control transfers from step 610 to step 612.

[0023] In step 612 a new frame is initiated. From step 612, control transfers to step 614 where the routine 600 stores any outstanding isochronous requests from a frame immediately prior to the new frame. One skilled in the art will appreciate these requests might be stored in a variety of ways (e.g., FIFO buffer or table). From step 614, control transfers to step 616. In step 616, the routine 600 processes any isochronous requests from two frames prior to the new frame in a round-robin fashion at the start of the new frame. From step 616, control transfers to step 604. From step 604, if a bus request is not for an asynchronous transaction, then control transfers to 606. In step 606, the routine 600 processes isochronous requests and removes any serviced isochronous transaction from the list created in step 614. From step 606, control returns to step 604.

[0024] Thus, with this technique, an isochronous transaction is delayed, at most, three frames before being processed. That is, an isochronous request at the beginning of one frame is serviced by the end of a second successive frame. As in Figure 5, asynchronous transactions in a current frame are processed with priority over an isochronous transaction of the current frame. Further, asynchronous transactions in a current frame are processed with priority over an isochronous transaction of a frame immediately preceding the current frame. One skilled in the art will appreciate the sort of logic or code helpful to implement the techniques described in connection with Figures 5 and 6.

[0025] Turning to Figure 7, an embodiment of the present invention which is implemented within a network interface card (NIC) 128 is illustrated. When the shared bus is an Ethernet, the type of frame can be determined through examining an address field (sending or receiving) or control field of the frame. The NIC 128 is included within the computer system S and enables the system S to transmit and receive data across a shared bus 700. A typical network driver routine implements code to allow a device to drive information onto the shared bus 700. The driver routine may reside in firmware 702 on the network interface

card 128 or be stored elsewhere within the computer system S. The device could be of a variety of types, such as a server or a workstation (represented by the computer system S). The network driver routine causes the information that is to be transferred to be driven onto the shared bus 700. Figure 7 illustrates the NIC 128 as containing the network driver routine (firmware) 702 which schedules asynchronous and isochronous transactions of a particular device on the shared bus 700 according to the present invention. It should be understood that the disclosed scheduling techniques can be implemented in hardware or software.

[0026] Thus, asynchronous and isochronous transactions on a shared bus are scheduled such that asynchronous latency is generally reduced while maintaining isochronous latency below an acceptable level. When a bus request is received the technique determines whether the bus request in a current frame is for an asynchronous transaction or an isochronous transaction. If an asynchronous bus request has been initiated it is processed, otherwise an isochronous bus request is processed. Bus requests for an isochronous transaction are queued if received while an asynchronous transaction is currently being processed. Asynchronous transactions are given priority until a current frame time has ended. In one disclosed embodiment, at the start of a new frame any queued isochronous transactions are processed before asynchronous transactions of the current frame are given priority. In another disclosed embodiment, queued isochronous transactions are only processed at the start of the new frame if they are from two frames prior to the new frame. Either approach normally optimizes bandwidth between asynchronous and isochronous traffic.

[0027] The foregoing disclosure and description of the various embodiments are illustrative and explanatory thereof, and various changes in the components, circuit elements, bus characteristics, data characteristics, code, wiring connections and contacts, as well as in the details of the illustrated hardware and software and construction and method of operation may be made without departing from the spirit of the invention.